arXiv:1610.01417v1 [stat.ML] 5 Oct 2016

Decentralized Topic Modelling with Latent Dirichlet Allocation

Igor Colin *
LTCI, Télécom ParisTech
Université Paris-Saclay, 75013 Paris

Christophe Dupuy *
Technicolor; INRIA - Sierra Project - Team
2, rue Simone Iff, 75012 Paris

Abstract

Privacy preserving networks can be modelled as decentralized networks (*e.g.,* sensors, connected objects, smartphones), where communication between nodes of the network is not controlled by a master or central node. For this type of networks, the main issue is to gather/learn global information on the network (*e.g.,* by optimizing a global cost function) while keeping the (sensitive) information at each node. In this work, we focus on text information that agents do not want to share (*e.g.,* , text messages, emails, confidential reports). We use recent advances on decentralized optimization and topic models to infer topics from a graph with limited communication. We propose a method to adapt latent Dirichlet allocation (LDA) model to decentralized optimization and show on synthetic data that we still recover similar parameters and similar performance at each node than with stochastic methods accessing to the whole information in the graph.

1 Introduction

Decentralized networks, *i.e.,* networks with limited communication between nodes, provide an ideal framework for privacy preserving optimization. They are particularly adapted to the case where agents seek to minimize a global cost function which is separable in the data collected locally by each agent. In such networks, it is typically impossible to efficiently centralize data or to globally aggregate intermediate results: agents can only communicate with their immediate neighbors, often in a completely asynchronous fashion. In a privacy setting, agents never share the raw information they have but rather communicate aggregated information or parameters. Recent work on decentralized optimization focuses on convex optimization: such methods are mostly based on gradient descent [1, 2] or dual averaging [3, 4] and theoretical convergence upper bounds have been established for both synchronous and fully asynchronous computations.

In this paper, we tackle the non-convex problem of topic modelling, where agents have sensitive text data at their disposal that they can not or do not want to share (e.g., text messages, emails, confidential reports). More precisely, we adapt the particular Latent Dirichlet Allocation (LDA) [5] model to decentralized networks. We combine recent work of [6] on online inference for latent variable models, which adapts online EM [7] with local Gibbs sampling in the case of intractable latent variable models (such as LDA) and recent advances on decentralized optimization [3, 4]. The method presented in [6] is particularly adapted to decentralized framework as it consists in iteratively updating sufficient statistics, which can be done locally. After presenting our DELEDA (for Decentralized LDA) algorithm, we give a brief sketch of convergence proof. Then, we apply our new method to synthetic datasets and show that our method recovers the same parameters and has similar performance than the online method [6] after enough iterations.

*The two authors contributed equally.

30th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain.

2 Background

Latent Dirichlet allocation (LDA) [5]. Let D be the number of documents of a corpus $\mathcal{C} = \{X_1, \dots, X_D\}$, V the number of words in our vocabulary and K the number of latent topics in the corpus. Each topic β^k corresponds to a discrete distribution on the V words (that is an element of the simplex in V dimensions). A hidden discrete distribution θ_d over the K topics (that is an element of the simplex in K dimensions) is attached to each document d. LDA is a generative model applied to a corpus of text documents which assumes that each word of the d-th document X_d is generated as follows:

- Choose $\theta_d \sim \mathrm{Dirichlet}(\alpha)$,
- For each word $x_n \in X_i = \left(x_1, \dots, x_{N_{X_d}}\right)$:
  - Choose a topic $z_n \sim \mathrm{Multinomial}(\theta_d)$,
  - Choose a word $x_n \sim \mathrm{Multinomial}(\beta^{z_n})$.

In LDA, the inference consists in learning the topic matrix $\beta \in \mathbb{R}^{K \times V}$ — global parameter — and topic proportions θ_d of each document — local variable — from a corpus of documents. We choose to adapt Gibbs Online EM for LDA (G-OEM) [6] to the decentralized framework. The other inference methods for LDA (*e.g.,* [5, 8, 9, 10, 11]) are very close to G-OEM but are less robust in practice (see [6] for complete analysis).

Gibbs Online EM for LDA [6]. We consider an exponential family model on random variables (X, h) with parameter $\eta \in \mathcal{E} \subseteq \mathbb{R}^d$ and with density [12]:

$$p(X, h|\eta) = a(X, h) \exp\left[\langle \phi(\eta), S(X, h)\rangle - \psi(\eta)\right], \tag{1}$$

with h hidden and X observed. G-OEM is an adaptation of the online EM algorithm [7] that maximizes the likelihood of observed data $\max_{\eta \in \mathcal{E}} \sum_{i=1}^{N} \log p(X_i|\eta)$ in the case where the conditional $p(h|X, \eta)$ is intractable to compute, such as for LDA. In particular, the G-OEM algorithm iteratively updates the sufficient statistics s of the model through the formula at step $t + 1$:

$$s^{t+1} = (1 - \rho_{t+1})s^t + \rho_{t+1}\mathbb{E}_{p(h_{t+1}|X_{t+1}, \eta^*(s^t))}[S(X_{t+1}, h_{t+1})], \tag{2}$$

where $\eta^*(s)$ is the maximum of likelihood for the sufficient statistics s, *i.e.,* $\eta^*(s) \in \arg\max \langle \phi(\eta), s\rangle - \psi(\eta)$ (which is the usual M-step update in the EM algorithm [13]). For LDA, the expectation in (2) is intractable to compute and is approximated with Gibbs sampling — see [6] for more details.

In this setting for LDA, an observation is a document X_i. Each observation X_i is associated with the hidden variables h_i, with $h_i \equiv (Z_i = (z_1, \dots, z_{N_{X_i}}), \theta_i)$. The vector θ_i represents the topic proportions of document X_i and Z_i is the vector of topic assignments of each word of X_i. The hidden variable h_i is local, *i.e.,* attached to one observation X_i. The parameters of the model are global, represented by $\eta \equiv (\beta, \alpha)$, where β represents the topic matrix and α represents the Dirichlet prior on topic proportions.

Decentralized setting Let $n > 0$ and let $\mathcal{G} = ([n], \mathcal{E})$ be an undirected graph of n agents. We consider a setting where each agent aims to optimize a global ojective, without revealing its local observations. In addition, we are interested in the case where communication is possibly asynchronous, *i.e.,* there is no global clock ensuring iterations synchrony between nodes. Gossip algorithms [14] are particularly adapted to this setting, as they rely on peer-to-peer interactions. Initially, each agent $i \in [n]$ stores an iterates s_i^0; at each iteration, a node awakes and select one of its neighbor: the two nodes then average their iterates. Then, in the synchronous (resp. asynchronous) setting, all the nodes (resp. the two active nodes) update their iterate according to a local rule, say (2).

When performing decentralized *estimation e.g.,* estimating the sample mean, the averaging step allows for a quick uniformization of the network iterates. Indeed, nodes iterates converge to the network average at an exponential rate which depends on some communication property of the graph called the spectral gap: the larger the spectral gap, the faster the convergence. In the case of decentralized *optimization*, since the network is not conserving its mass during the process — typically, gradients are added —, the convergence is not guaranteed in general. However, provided that at each iteration the mass added to the network is bounded, one can formulate upper bounds on the convergence rate, usually polynomial in the number of iterations.

Algorithm 1 Decentralized LDA in the synchronous setting.

Require: Step size $(\rho_t)_{t\geq 1} > 0$.
1: Each node $i \in [n]$ initializes s_i at s_i^0.
2: **for** $t = 1, \dots, T$ **do**
3: Draw (i, j) uniformly at random from $\mathcal{E}$
4: Set $s_i, s_j \leftarrow \frac{s_i + s_j}{2}$
5: **for** $k = 1, \dots, n$ **do**
6: Update (locally) sufficient statistic s_k with 2
7: **end for**
8: **end for**
9: **return** Each node k has s_k

3 DELEDA algorithm

In this paper, we aim at combining both G-OEM for LDA and decentralized optimization. Since the G-OEM updates are very similar to stochastic gradient updates, the reasoning of decentralized convex optimization can be extended to G-OEM. In a synchronous setting, our algorithm coined DELEDA consists in the following: each node $i \in [n]$ initializes its sufficient statistics to s_i^0. At each iteration, the two awaken nodes average their statistics. Then, every node performs a local update of its statistics using update rule (2) and its own observations. The method is detailed in Algorithm 1.

For $t \geq 0$, let us define the average iterate $\overline{s}^t = n^{-1} \sum_{i=1}^n s_i^t$. Using Algorithm 1 description and update rule (2), one can see that $\overline{s}^{t+1}$ is obtained from $\overline{s}^t$ using the G-OEM update rule with all observations of the network. Therefore, $\overline{s}^t$ convergence derives from G-OEM theoretical guarantees [6]; in order to prove the convergence of each agent, one only has to show that for all $i \in [n]$, s_i^t tends to $\overline{s}^t$ as the number of iterations t grows.

Let us define the iterates matrix $\mathbf{S}^t = (s_1^t, \dots, s_n^t)^\top$ and the averaging operator $\mathbf{W}^t = \mathbf{I}_n - \frac{1}{2}(e_i - e_j)(e_i - e_j)^\top$, where i and j are the nodes awaken at iteration t. An iteration of Algorithm 1 can be reformulated as follows:

$$\mathbf{S}^{t+1} = (1 - \rho_{t+1})\mathbf{W}^{t+1}\mathbf{S}^t + \rho_{t+1}\mathbf{G}^{t+1},$$

where $\mathbf{G}^{t+1}$ is the matrix of expectations in (2). Recursively, one can obtain:

$$\mathbf{S}^t = \sum_{r=0}^{t} (1 - \rho_t) \dots (1 - \rho_{r+1}) \rho_r \mathbf{W}^t \dots \mathbf{W}^r \mathbf{G}^r.$$

Using properties of averaging matrices[2], one can rearrange the above expression and write:

$$\begin{aligned}
\mathbb{E}\|\mathbf{S}^t - \overline{s}^t \mathbf{1}_n^\top\| &\leq \sum_{r=0}^{t} (1 - \rho_t) \dots (1 - \rho_{r+1}) \rho_r \mathbb{E}\left\|(\mathbf{W}^t - \mathbf{1}_{n\times n}) \dots (\mathbf{W}^r - \mathbf{1}_{n\times n})\mathbf{G}^r\right\| \\
&\leq \sum_{r=0}^{t} \rho_r \lambda_2^{\frac{t-r}{2}} \mathbb{E}\|\mathbf{G}^r\|, \qquad (3)
\end{aligned}$$

where λ_2 is the second largest eigenvalue of $\mathbb{E}[\mathbf{W}^\mathbf{t}]$. If the graph $\mathcal{G}$ is connected and non-bipartite, one can show that $0 < \lambda_2 < 1$. Therefore, when $(\rho_t)_{t\geq 0}$ tends towards 0 and provided that $\mathbb{E}\|\mathbf{G}_r\|$ is bounded, the expected gap between nodes iterates and the average $\overline{s}^t$ goes to 0. In practice, this last assumption holds for instance when the Dirichlet parameter α is either fixed or lower-bounded, *i.e.,* $\alpha > r > 0$. Also, as evidenced by inequality (3), one should expect DELEDA to converge faster on a graph with high spectral gap, as it is proportional to $1 - \lambda_2$ [4, 14].

Remark 1. Note that this reasoning holds in the synchronous setting. In the fully asynchronous setting, one needs to take into account the nodes degrees and adapt the update rule (2) for the algorithm to converge to the true objective, as detailed for instance in [4].

[2]See [14, 4] for more detailed analysis of averaging matrices.

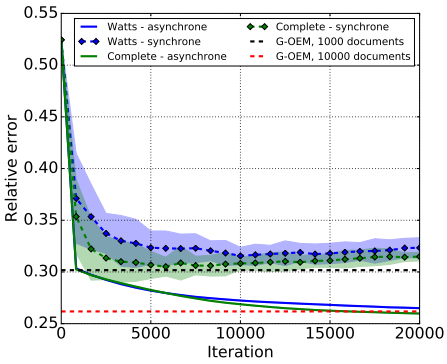

(a) Relative error $(LP - LP^*)$.

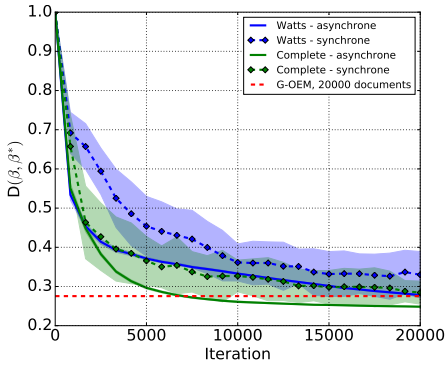

(b) Average distance to the optimal parameter β^*.

Figure 1: Distance of log-perplexity (a) and parameters (b) as a function of the iterations.

4 Numerical experiments

Dataset. We perform experiments on a synthetic graph of n=50 nodes. Given these nodes, we consider the complete graph ($|\mathcal{E}| = 1225$ edges) and a graph where the edges are generated through a Watts-Strogatz model [15], with 100 edges and rewiring probability $p = 0.3$. We attach 20 documents at each node, generated through LDA generative process, with a vocabulary size $V = 100$ and $K = 5$ topics. The length of each document is drawn from a Poisson with parameter 10. We compare our algorithm to — centralized — G-OEM applied to the 1000 synthetic documents of the graph with a batch size of 20 documents. In particular, at each iteration of the G-OEM algorithm we choose uniformly at random 20 documents among the 1000 of the total dataset, and update η according to the G-OEM process. In practice, we update β at each iteration and let $\alpha = \alpha^*$ fixed, as often done in previous work [8, 11, 16, 17].

Evaluation. We evaluate our method by computing the likelihood on held-out documents, that is $p(X|\eta)$ for any test document X. For LDA, the likelihood is intractable to compute. We approximate $p(X|\eta)$ with the "left-to-right" evaluation algorithm [18] applied to each test document. In the following, we present results in terms of log-perplexity, defined as the opposite of the log-likelihood $LP = -\log p(X|\eta)$. The lower the log-perplexity, the better the corresponding model. In our experiments, we compute the average test log-perplexity over test documents. More precisely, we display the relative error $(LP/LP^* - 1)$, where LP^* is the average log-perplexity of test documents computed with optimal parameters η^* that generated the data.

We also compare the performance in terms of distance to the optimal topic matrix β^*, which generated the documents. As the performance is invariant to topic permutation in LDA, we compute the following distance between the infered parameter β and optimal parameter β^*:

$$D(\beta, \beta^*) = \min_{M \in \mathbb{R}^{K \times K}} \frac{\|M\beta - \beta^*\|_F}{\|\beta^*\|_F} = \frac{\|\beta^*\beta^\top(\beta\beta^\top)^{-1} - \beta^*\|_F}{\|\beta^*\|_F},$$

which is invariant by topic permutation (*i.e.,* by permutation of rows of β).

Results for asynchrone and synchrone method are presented in Figure 1. We observe that our method converges to the same state than G-OEM after enough iterations. In particular, each agent is able to recover the topic matrix that generated all the documents without any direct access to documents of the other nodes.

As expected, DeLeDA converges faster on a perfectly connected network, even though the "well-connected" design of the Watts-Strogatz network allows it to stay close to the complete graph, perplexity-wise. Finally, the synchronous version of DeLeDA appears to converge slower than the asynchronous one. This is an overfitting issue: synchronous computations lead each node to perform roughly $n/2$ local updates between two averaging steps, whereas asynchronous computations force nodes to update only once between two averaging steps.

References

[1] Angelia Nedic and Asuman Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61, 2009.

[2] Angelia Nedić. Asynchronous broadcast-based convex optimization over a network. *Automatic Control, IEEE Transactions on*, 56(6):1337–1351, 2011.

[3] Alekh Agarwal, Martin J Wainwright, and John C Duchi. Distributed dual averaging in networks. In *Advances in Neural Information Processing Systems*, pages 550–558, 2010.

[4] I. Colin, A. Bellet, J. Salmon, and S. Clémençon. Gossip dual averaging for decentralized optimization of pairwise functions. *Proc. ICML*, 2016.

[5] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent Dirichlet allocation. *JMLR*, 3:993–1022, 2003.

[6] C. Dupuy and F. Bach. Online but accurate inference for latent variable models with local gibbs sampling. *arXiv preprint arXiv:1603.02644*, 2016.

[7] O. Cappé and E. Moulines. Online EM algorithm for latent data models. *Journal of the Royal Statistical Society*, 71(3):593–613, 2009.

[8] M. D. Hoffman, D. M. Blei, and F. R. Bach. Online learning for latent Dirichlet allocation. *Adv. NIPS*, 2010.

[9] T. L. Griffiths and M. Steyvers. Finding scientific topics. *Proceedings of the National Academy of Sciences*, 101(suppl 1):5228–5235, 2004.

[10] I. Sato, K. Kurihara, and H. Nakagawa. Deterministic single-pass algorithm for LDA. *Adv. NIPS*, 2010.

[11] D. Mimno, M. Hoffman, and D. Blei. Sparse stochastic inference for latent Dirichlet allocation. *Proc. ICML*, 2012.

[12] E. L. Lehmann and G. Casella. *Theory of point estimation*, volume 31. Springer Science & Business Media, 1998.

[13] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the royal statistical society. Series B (methodological)*, 39(1):1–38, 1977.

[14] Stephen Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. Randomized gossip algorithms. *IEEE Trans. Inf. Theory*, 52(6):2508–2530, 2006.

[15] Duncan J Watts and Steven H Strogatz. Collective dynamics of 'small-world'networks. *Nature*, 393(6684):440–442, 1998.

[16] S. Patterson and Y. W. Teh. Stochastic gradient Riemannian langevin dynamics on the probability simplex. *Adv. NIPS*, 2013.

[17] Y. Gao, J. Chen, and J. Zhu. Streaming gibbs sampling for LDA model. *arXiv preprint arXiv:1601.01142*, 2016.

[18] H. M. Wallach, I. Murray, R. Salakhutdinov, and D. Mimno. Evaluation methods for topic models. *Proc. ICML*, 2009.